EECS C145B / BioE C165: Image Processing and Reconstruction Tomography

Lecture 3

Jonathan S. Maltz

jon@eecs.berkeley.edu http://muti.lbl.gov/145b
510-486-6744

Topics to be covered

1. The inverse problem of deconvolution (image restoration)
2. Motivation for studying the discrete Fourier transform in 2+D
3. Review of the 1D Fourier transform
4. The discrete Fourier transform
5. Introduction to the 2D FT
6. The fast Fourier transform (self study)

Reading

1. Gonzalez and Woods pp. 147-157.

Self-study: The fast Fourier transform

1. Gonzalez and Woods pp. 208-213.

Refresher material

1. Continuous time Fourier transform: Oppenheim and Willsky with Young (1983), pp. 190-214, 223-225
2. Discrete time Fourier transform: Oppenheim & Schafer, pp. 40-65.
3. Sampling: Oppenheim & Schafer, pp. 140-153.

2D linear shift invariant systems

$$f[m,n] \longrightarrow \boxed{h[m,n]} \longrightarrow g[m,n]$$

Inverse problem: We have an image **g** and know what **h** is and want to find **f**. This is **image restoration**.

Example I: Correcting myopia (short sightedness)

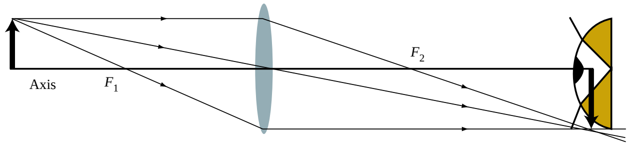

The optometrist measures **h** using instruments and makes a lens which modifies the combined point spead function (PSF) so that the image is focussed on the retina.

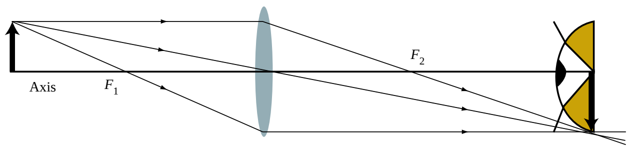

Exercise: Draw in the type of lens used.

Example II: Correcting the Hubble Space Telescope

The main mirror of the Hubble space telscope was incorrectly designed. This imaging system should have had a PSF as close to a single impulse as possible (i.e., no point spread). This was not the case and the first images were very disappointing.

Photo courtesy of the Space Telescope Science Institute

Example II: Correcting the Hubble Space Telescope

To make the PSF (**h**) better approximate an impulse, five space walks were conducted during a single shuttle flight. Corrective mirrors were added to make the Hubble images crisp. The corrective system's transfer function was designed to be the inverse of that of the mirror defect.

Photo courtesy of the Space Telescope Science Institute

Digital image restoration

Assume $h[m,n]$ is the blurring PSF:

$$f[m,n] \longrightarrow \boxed{h[m,n]} \longrightarrow g[m,n]$$

We try to find an approximation to the inverse of this function $\hat{h}^{-1}[m,n]$. If we feed the blurred image through the inverse system, we get $\hat{f}[m,n]$. This image is an approximation of the true image $f[m,n]$.

$$g[m,n] \longrightarrow \boxed{\hat{h}^{-1}[m,n]} \longrightarrow \hat{f}[m,n]$$

This inverse operation is called **deconvolution**.

Hubble example continued

Right now, inside the Hubble telescope, the corrective mirrors are convolving the blurred image $g(x, y)$ with a PSF $\hat{h}^{-1}(x, y)$ to produce $\hat{f}(x, y)$, a good approximation to the desired image $f(x, y)$.

2D deconvolution as a matrix operation

Recall the example from Lecture 2:

$$\mathsf{g} = \begin{bmatrix} \mathbf{g}_0 \\ \mathbf{g}_1 \\ \mathbf{g}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{H}_0 & \mathbf{0} \\ \mathbf{H}_1 & \mathbf{H}_0 \\ \mathbf{0} & \mathbf{H}_1 \end{bmatrix} \triangleq \mathcal{H}\mathsf{f}$$

Deconvolution can be performed if we can find a suitable approximation to the inverse of $\mathcal{H}$:

$$\hat{\mathsf{f}} = \hat{\mathcal{H}}^{-1}\mathsf{g}$$

Since this is equivalent to solving a set of simultaneous linear equations, why can't we just use the normal matrix inverse $\mathcal{H}^{-1}$?

____

____

If $f[m, n]$ is a 512×512 image, estimate the size of $\mathcal{H}$ __________

How much memory do we need to store $\mathcal{H}$ if one double precision floating point number requires 8 bytes? __________

Motivation for studying the mutidimensional FT

1. As is the case for signals in 1D, the frequency spectra of images reveal important properties of images and the systems that created them.
2. Many topics in this course center around the linear inverse problem:

$$\hat{\mathsf{f}} = \hat{\mathcal{H}}^{-1}\mathsf{g}$$

   We will find that the FT (via the convolution theorem) enables us to perform deconvolution and filtering of large images in a computationally feasible way.

The discrete Fourier transform (DFT)

Forward transform (analysis):

$$G[k] = \begin{cases} \sum_{n=0}^{N-1} g[n] \mathrm{e}^{-\jmath 2\pi kn/N}, & 0 \le k \le N-1, \\ 0 & \text{otherwise} \end{cases}$$

Inverse transform (synthesis):

$$g[n] = \begin{cases} \frac{1}{N} \sum_{k=0}^{N-1} G[k]\, \mathrm{e}^{\jmath 2\pi kn/N}, & 0 \le n \le N-1, \\ 0 & \text{otherwise} \end{cases}$$

Fourier transform: There is only one.

It is called different names, and different notation is used, depending on the nature of the signal to which it is applied.

Real space axis	Real space function	Frequency axis	Spectrum	Type of FT
continuous	aperiodic	continuous	aperiodic	Fourier transform
continuous	periodic	continuous	discrete	Fourier series
discrete	aperiodic	continuous	countinuous, periodic	Discrete space Fourier transform (DSFT)
discrete	**periodic**	**discrete**	**periodic**	Discrete Fourier transform (DFT)

Note: Only the DFT is suitable for direct implementation on a computer because both the input and output are discrete arrays. **However, any signal analyzed by the DFT is implicitly assumed to be periodic.**

Forward 1D DFT: Example

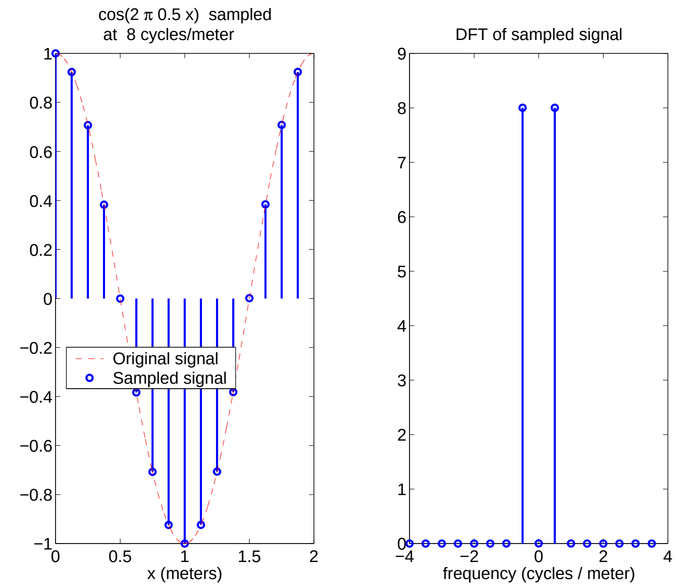

DFT assumes we are giving it a single period of a periodic function

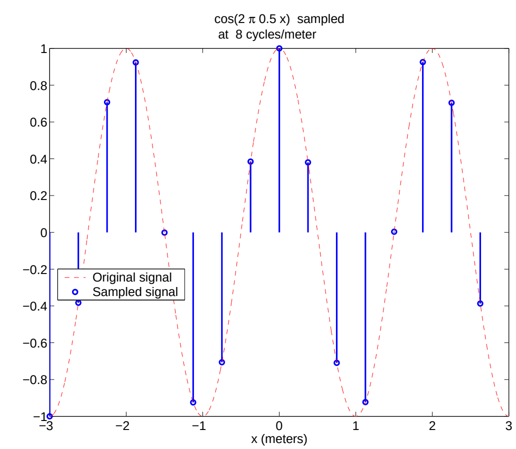

The signal shown here is the signal that is actually analyzed by the DFT.

Forward 1D DFT: Example

The 1D DFT takes a vector on size N (16 in this case) as input:

$$\text{Samples of } \cos(2\pi 2t): \begin{bmatrix} 1.00 \\ 0.92 \\ 0.71 \\ 0.38 \\ -0.00 \\ -0.38 \\ -0.71 \\ -0.92 \\ -1.00 \\ -0.92 \\ -0.71 \\ -0.38 \\ 0.00 \\ 0.38 \\ 0.71 \\ 0.92 \end{bmatrix}$$

Forward 1D DFT: Example

And returns a vector on size N (16 in this case) as output:

$$\text{DFT of vector of samples:} \begin{bmatrix} 0.00 \\ 8.00 + 0.01i \\ -0.00 - 0.00i \\ -0.00 - 0.00i \\ -0.00 - 0.00i \\ -0.00 - 0.00i \\ -0.00 - 0.00i \\ -0.00 - 0.00i \\ -0.00 \\ -0.00 + 0.00i \\ -0.00 + 0.00i \\ -0.00 + 0.00i \\ -0.00 + 0.00i \\ -0.00 + 0.00i \\ -0.00 + 0.00i \\ 8.00 - 0.01i \end{bmatrix}$$

Axis of the DFT

We sampled the continuous cosine signal 8 times over 1 meter. Thus our sampling frequency is:

$$f_s = 8 \text{ cycles / m}$$

The input vector had $N = 16$ samples. This means that the output vector will also have $N = 16$ samples.
To what frequency does each element correspond?

Axis of the DFT

Element number	frequency (cycles/m)	
0	$0 \times f_s/N = 0 \times 8/16 =$	0
1	$1 \times f_s/N = 1 \times 8/16 =$	0.5
2	$2 \times f_s/N = 2 \times 8/16 =$	1
$\vdots$	$\vdots$	
7	$7 \times f_s/N = 7 \times 8/16 =$	3.5
8	$8 \times f_s/N - f_s = 8 \times 8/16 - 8 =$	−4
9	$9 \times f_s/N - f_s = 9 \times 8/16 - 8 =$	−3.5
$\vdots$	$\vdots$	
14	$14 \times f_s/N - f_s = 14 \times 8/16 - 8 =$	−1
15	$15 \times f_s/N - f_s = 15 \times 8/16 - 8 =$	−0.5

DFT algorithms return a specific single period of the DFT

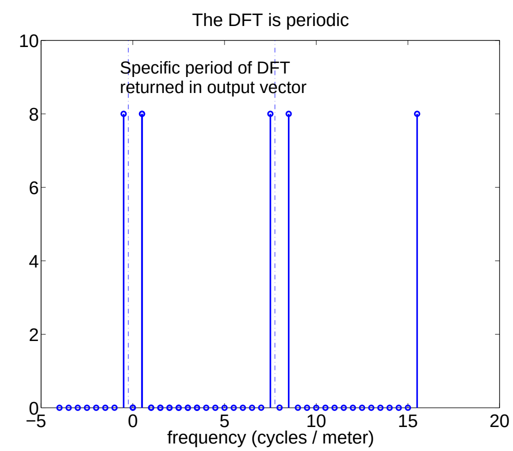

Frequency spectra are normally viewed with zero frequency at the center of the plot. This is why we place the DFT samples for frequencies $\geq fs$ to the left of sample 0.

Forward 1D DFT: Example of DFT: Non-integral number of periods

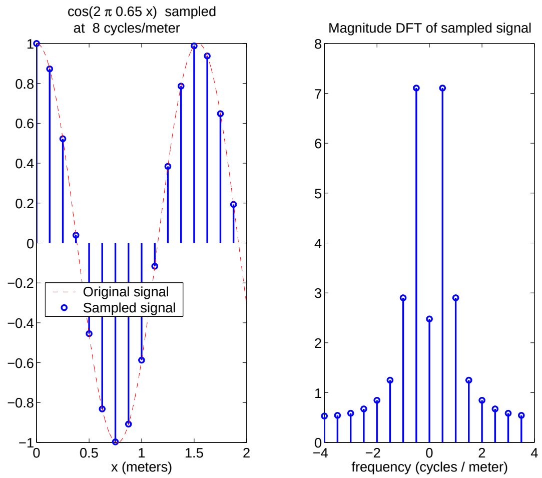

Why do we have so many non-zero peaks in the spectrum?

Answer: DFT assumes we want to transform this:

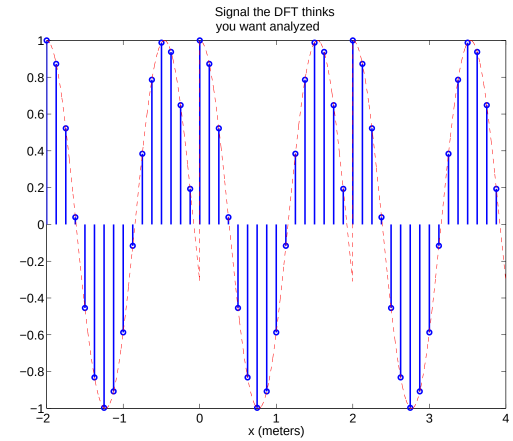

The implicit discontinuity introduces other components into the spectrum.

Where does this periodicity come from?

Review 1D continuous space Fourier transform

Forward transform: (analysis)

$$F(\omega) = \int_{-\infty}^{\infty} f(x)\, \mathrm{e}^{-\jmath\, \omega x} dx$$

Inverse transform: (synthesis)

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} F(\omega)\, \mathrm{e}^{\jmath\, \omega x} d\omega$$

An continuous space signal with limited support has a continuous FT with infinite support.

Example:

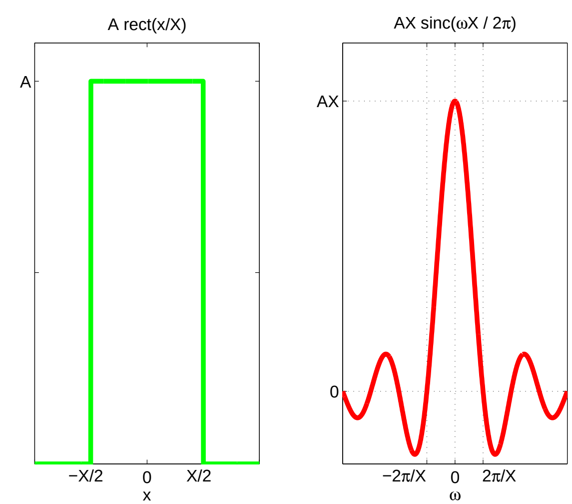

sinc and rect

Recall:

$$\mathrm{sinc}(x) \triangleq \frac{\sin(\pi x)}{\pi x}$$

$$\mathrm{rect}\left(\frac{x}{X}\right) = \begin{cases} 1, & |x| < X/2 \\ 0, & |x| > X/2 \end{cases}$$

A continuous signal that is periodic in space has a discrete FT with finite support.

Example:

Duality and sampling

Duality: If we have a Fourier transform pair:

$$g(x) \rightleftharpoons f(\omega)$$

Then we may find the transform of $f(x)$ as follows:

$$f(x) \rightleftharpoons 2\pi g(-\omega)$$

Because of the duality between real space and Fourier space, if a **periodic real space signal** has a **discrete FT**, a **discrete real space signal** must have a **periodic FT**. Consequently, if we sample a signal, we expect its FT to be periodic.

Recall

1. A comb of impulses in real space separated by a distance X transforms to a comb of impulses in frequency space separated by $2\pi/X$:

$$\sum_{k=-\infty}^{\infty} \delta(x - kX) \rightleftharpoons \sum_{k=-\infty}^{\infty} \delta(\omega - k\,2\pi/X)$$

2. The convolution property tells us that:

$$f(x) * h(x) \rightleftharpoons F(\omega)H(\omega)$$

3. The modulation property:

$$f(x)h(x) \rightleftharpoons \frac{1}{2\pi}\Big(F(\omega) * H(\omega)\Big)$$

Note how (2) can be derived from (3) and vice versa using the duality property.

Sampling a continuous space signal

Consider sampling a signal $f(x)$ at a sampling interval of X:

$$f_s(x) = \sum_{k=-\infty}^{\infty} \delta(x - kX) f(x)$$

Taking the Fourier transform (using the convolution property) gives:

$$F_s(\omega) = \frac{2\pi}{X} \sum_{k=-\infty}^{\infty} \delta(\omega - k\,2\pi/X) * F(\omega)$$

We may now use the “rubber stamp” method to visualize the spectrum of $F_s(\omega)$.

Example spectrum of a sampled signal

After sampling

In the real world, continuous signals or images are converted to arrays of numbers using an analog-to-digital (A/D) converter.

- When we plot a sample of a signal on a continuous axis, we represent the sample using a Dirac delta distribution.
- When we plot a signal on a discrete space, we represent the sample using a Kronecker delta (unit sample).

Representation of samples

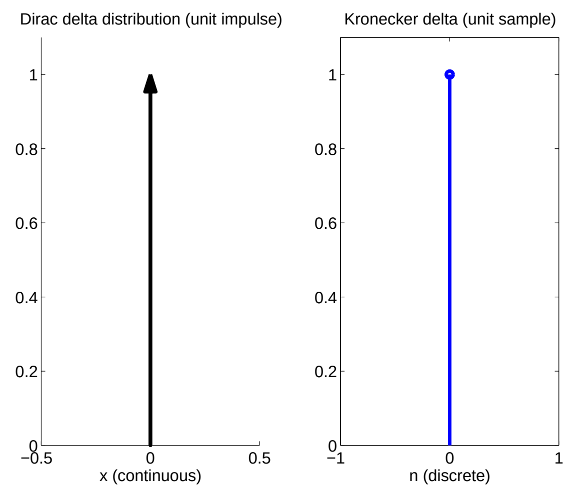

The discrete space Fourier transform (DSFT)

The DSFT is the Fourier transform of a sampled signal. It is, in general, continuous. The DSFT of a signal is given by:

$$F(e^{\jmath\omega}) = \sum_{n=-\infty}^{\infty} f[n] e^{-\jmath\omega n}$$

1. The DSFT axis is normalized in frequency. If a continuous signal is sampled at a sampling frequency $\omega = 2\pi/X$, this frequency is mapped to 2π on the $e^{\jmath\omega}$ axis.
2. We saw earlier that the FT of a sampled signal is periodic and repeats every $2\pi/X$. The DSFT is normalized with respect to sampling frequncy and has a period of 2π.

DSFT transform pair example

$$\frac{\sin \omega_c n}{\pi n} \rightleftharpoons X(e^{\jmath\omega}) = \begin{cases} 1, & |\omega| < \omega_c \\ 0, & \omega_c < |\omega| \leq \pi \end{cases}$$

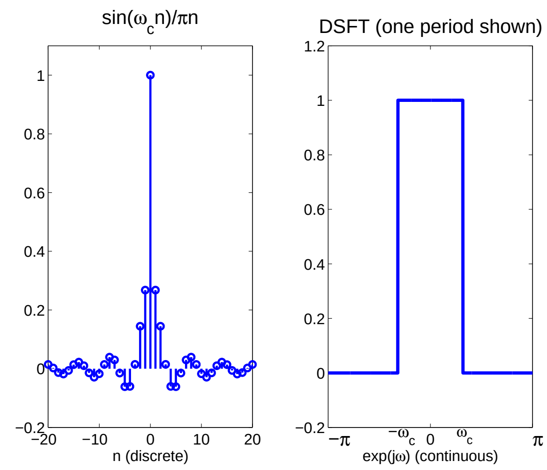

DFT: Sampling the DSFT

- The DFT coefficients $G[k]$ are samples of a single period of the DSFT. We know, however, that when we sample a signal in one domain, its transform in the other domain is periodic.
- The DFT provides a practical method of finding the samples of the DSFT from the samples of a signal and vice versa.
- We "pay" for this convenience by having to accept the assumption of periodicity.

DFTs of aperiodic signals: Windowing

- Whenever possible, try to take a DFT of an entire period of the input signal.
- If the input signal is aperiodic, choose a segment of the signal such that the first value and last value are similar. This reduces edge discontinuities.
- Often, the input signal is multiplied by a window function that slowly tapers the edges of the signal down towards zero.

Common windowing functions

rectangular (implicitly assumed by DFT):

$$w[n] = \begin{cases} 1, & 0 \le n \le N \\ 0, & \text{otherwise} \end{cases}$$

von Hahn window:

$$w[n] = \begin{cases} (0.5 - 0.5\cos(2\pi n/N)), & 0 \le n \le N \\ 0, & \text{otherwise} \end{cases}$$

Hamming window:

$$w[n] = \begin{cases} (0.54 - 0.46\cos(2\pi n/N)), & 0 \le n \le N \\ 0, & \text{otherwise} \end{cases}$$

Plots of windowing functions

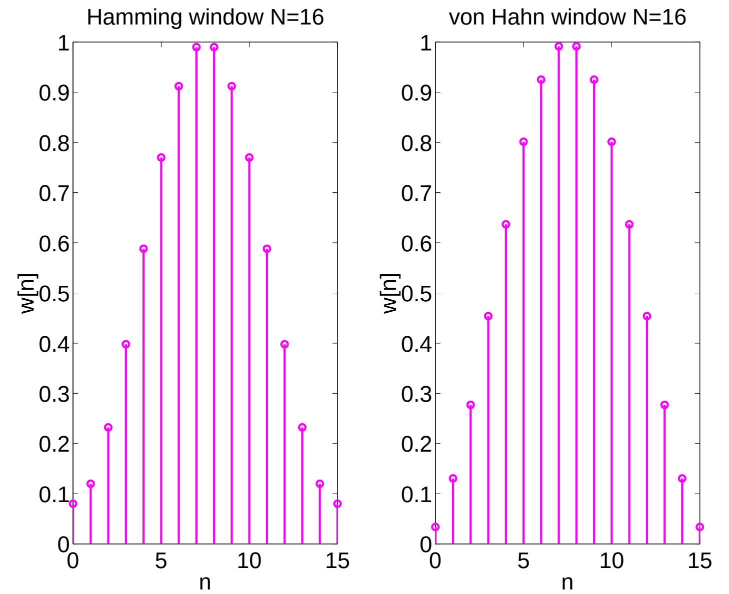

Windowed DFT: Example of DFT: Non-integral number of periods

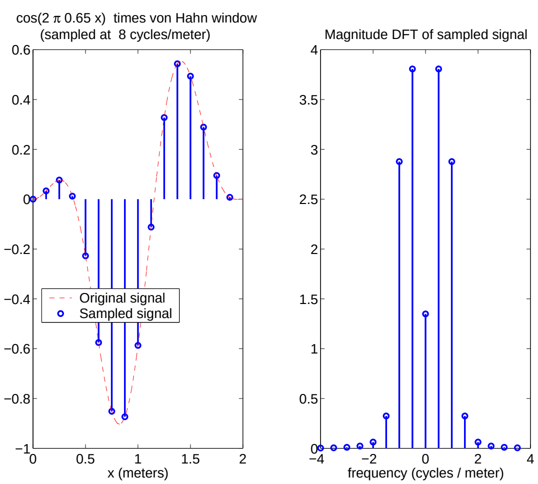

Windowing reduces the spectral contamination due to edge discontinuities

This is what the spectrum looked like before:

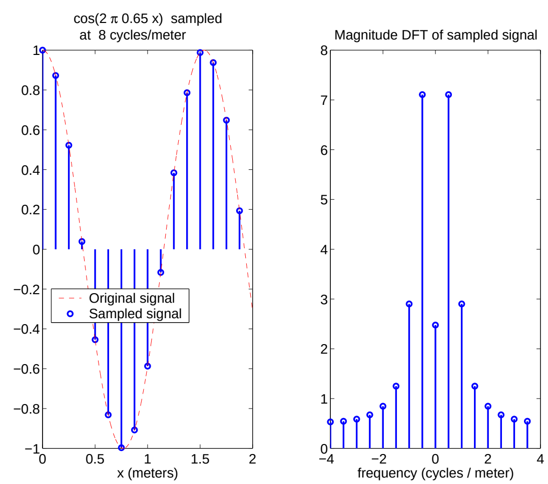

This time the DFT transformed this:

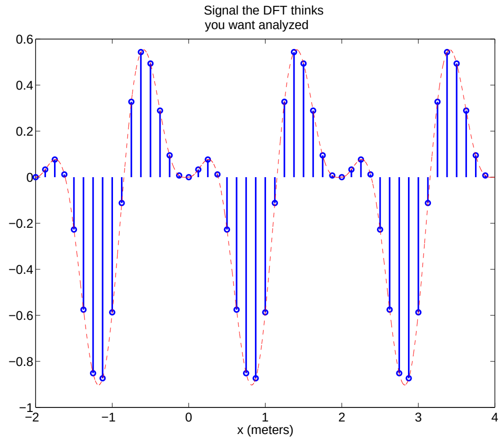

With this signal (because it **is** periodic), a better approach would have been to choose an integral number of periods and use a rectangular window.

The 2D continuous space Fourier transform

Forward transform: (analysis)

$$F(w_1, w_2) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x_1, x_2)\, \mathrm{e}^{-\jmath(w_1 x_1 + w_2 x_2)} dx_1 dx_2$$

Inverse transform: (synthesis)

$$f(x_1, x_2) = \frac{1}{2\pi} \int_{-\infty}^{\infty} F(\omega_1, \omega_2)\, \mathrm{e}^{\jmath(\omega_1 x_1 + \omega_2 x_2)} d\omega_1 d\omega_2$$

The multidimensional continuous space Fourier transform

Forward transform: (analysis)

$$F(\boldsymbol{\omega}) = \int_{\Re^n} f(\mathbf{x})\, \mathrm{e}^{-\jmath\, \boldsymbol{\omega} \cdot \mathbf{x}} d\mathbf{x}$$

Inverse transform: (synthesis)

$$f(\mathbf{x}) = (2\pi)^{-n} \int_{\Re^n} F(\boldsymbol{\omega})\, \mathrm{e}^{\jmath\, \boldsymbol{\omega} \cdot \mathbf{x}} d\boldsymbol{\omega}$$

The 2D continuous space Fourier transform: Notation issues

Most image processing texts give Fourier transform pairs for conventional rather than angular frequency. Gonzalez and Woods uses $x \triangleq x_1$, $y \triangleq x_2$, u for horizontal frequency and v for vertical frequency.

Forward transform: (analysis)

$$F(u, v) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y)\, \mathrm{e}^{-\jmath\, 2\pi(ux+vy)} dx dy$$

Inverse transform: (synthesis)

$$f(x, y) = \int_{-\infty}^{\infty} F(u, v)\, \mathrm{e}^{\jmath\, 2\pi(ux+vy)} dx dy$$

Example 2D FT pair

$$\cos(2\pi u_0 x + 2\pi v_0 y) \rightleftharpoons \frac{1}{2}\Big[\delta(u - u_0, v - v_0) + \delta(u + u_0, v + v_0)\Big]$$